

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 17-11-2015		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Nov-2011 - 31-May-2015	
4. TITLE AND SUBTITLE Final Report: Assessing Trustworthiness in Social Media: A Social Computing Approach			5a. CONTRACT NUMBER W911NF-11-1-0517		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Huan Liu			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Arizona State University ORSPA P.O. Box 876011 Tempe, AZ 85287 -6011			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 60361-LS.36		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Social media is gaining popularity in recent years and increasingly becoming an integral part of our life. Given the extensiveness, instantaneity, and diffusion speed of social media, e.g., a tweet or a clip of video, could galvanize a digital revolution or wreak havoc with one's otherwise routine and uneventful working life. With the presence of adversaries, the convenient use of and low barrier of social media brings about new challenges. How well we address these challenges can directly influence our ability to manage information and misinformation, and the future role of social media as a reliable communication mechanism. One such pressing challenge is to assess					
15. SUBJECT TERMS Social Networks, Social Media, Social Media Mining, Trust, Information Provenance					
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT			c. THIS PAGE	Huan Liu
UU	UU	UU		19b. TELEPHONE NUMBER 480-727-7349	

Report Title

Final Report: Assessing Trustworthiness in Social Media: A Social Computing Approach

ABSTRACT

Social media is gaining popularity in recent years and increasingly becoming an integral part of our life. Given the extensiveness, instantaneity, and diffusion speed of social media, e.g., a tweet or a clip of video, could galvanize a digital revolution or wreak havoc with one's otherwise routine and uneventful working life. With the presence of adversaries, the convenient use of and low barrier of social media brings about new challenges. How well we address these challenges can directly influence our ability to manage information and misinformation, and the future role of social media as a reliable communication mechanism. One such pressing challenge is to assess information trustworthiness in social media. We propose to investigate research issues related to social media trustworthiness and its assessment by leveraging social research methods, developing new computational social methods, and creating novel approaches to social media data collection and sharing.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
08/19/2014 14.00	Yi Chang, Huan Liu, Jiliang Tang. Mining social media with social theories, ACM SIGKDD Explorations Newsletter, (06 2014): 0. doi: 10.1145/2641190.2641195
08/19/2014 15.00	Jiliang Tang, Xia Hu, Huan Liu. Social recommendation: a review, Social Network Analysis and Mining, (11 2013): 0. doi: 10.1007/s13278-013-0141-9
08/19/2014 20.00	Pritam Gundecha, Geoffrey Barbier, Huan Liu. User Vulnerability and its Reduction on a Social Networking Site, ACM Transactions on Knowledge Discovery from Data (TKDD), (12 2014): 0. doi:
11/17/2015 23.00	Jiliang Tang, Huiji Gao, Atish Das Sarma, Yingzhou Bi, Huan Liu. Trust Evolution: Modeling and Its Applications, IEEE Transactions on Knowledge and Data Engineering, (06 2015): 1724. doi: 10.1109/TKDE.2014.2382576
TOTAL:	4

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
08/06/2013 5.00	Jiliang Tang, Huiji Gao, Xia Hu, Huan Liu. Exploiting Homophily Effect for Trust Prediction, the 6th ACM International Conference on Web Search and Data Mining (WSDM 2013). 03-FEB-13, Rome, Italy. : ,
08/06/2013 11.00	Suhas Ranganath, Pritam Gundecha, Huan Liu. A Tool for Assisting Provenance Search in Social Media, the 22nd ACM International Conference on Information and Knowledge Management (CIKM 2013). 27-OCT-13, . : ,
08/06/2013 10.00	Zhuo Feng, Huan Liu , Pritam Gundecha. Seeking Provenance of Information in Social Media, the 22nd ACM International Conference on Information and Knowledge Management (CIKM 2013). 27-OCT-13, . : ,
08/06/2013 9.00	Jiliang Tang, Huiji Gao, Xia Hu, Huan Liu.. Context-Aware Review Helpfulness Rating Prediction, the ACM Conference Series on Recommender Systems. 12-OCT-13, . : ,
08/06/2013 8.00	Zhuo Feng, Pritam Gundecha, Huan Liu. Recovering Information Recipients in Social Media via Provenance, The IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013). 25-AUG-13, . : ,
08/06/2013 7.00	Pritam Gundecha, Suhas Ranganath, Zhuo Feng, Huan Liu. A Tool for Collecting Provenance Data in Social Media, the 19th ACM SIGKDD International Conference On Knowledge Discovery and Data Mining (SIGKDD 2013). 11-AUG-13, . : ,
08/06/2013 6.00	Jiliang Tang, Xia Hu, Huiji Gao, Huan Liu. Exploiting Local and Global Social Context for Recommendation, the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013). 03-AUG-13, . : ,
08/19/2014 16.00	Jiliang Tang, Xia Hu, Huan Liu. Is Distrust the Negation of Trust? The Value of Distrust in Social Media, ACM Hypertext conference. 01-SEP-14, . : ,
08/19/2014 17.00	Mohammad Ali Abbasi, Jiliang Tang, Huan Liu. Scalable Learning of Users' Preferences Using Networked Data, ACM Hypertext conference. 01-SEP-14, . : ,
08/19/2014 18.00	Xia Hu, Jiliang Tang, Huan Liu. Leveraging knowledge across media for spammer detection in microblogging, the 37th international ACM SIGIR conference. 05-JUL-14, Gold Coast, Queensland, Australia. : ,
08/19/2014 19.00	Xia Hu, Jiliang Tang, Huan Liu. Online Social Spammer Detection, AAAI Conference on Artificial Intelligence. 27-JUL-14, . : ,
08/22/2012 1.00	Jiliang Tang, Huiji Gao , Huan Liu, Atish Das Sarma. eTrust: Understanding Trust Evolution in an Online World, the Eighteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 13-AUG-12, . : ,
08/22/2012 2.00	Jiliang Tang, Huiji Gao, Huan Liu. mTrust: Discerning Multi-Faceted Trust in a Connected World, the fifth ACM international conference. 07-FEB-12, Seattle, Washington, USA. : ,

- 09/05/2012 3.00 Pritam Gundecha, Huan Liu. Mining Social Media : A Brief Introduction, the Institute for Operations Research and the Management Sciences (INFORMS) TutORials in Operations Research. 16-OCT-12, . . . ,
- 11/17/2015 26.00 Jiliang Tang, Huan Liu. Trust in Social Computing, 23rd International Conference on World Wide Web. 07-APR-14, . . . ,
- 11/17/2015 27.00 Jie Tang, Huan Liu, Jiliang Tang. Recommendation in social media, 20th ACM SIGKDD international conference. 23-AUG-14, New York, New York, USA. . . . ,
- 11/17/2015 28.00 Dongwon Lee, Huan Liu. LIKE and Recommendation in Social Media, 24th International Conference on World Wide Web Companion. 18-MAY-15, . . . ,
- 11/17/2015 35.00 Jiliang Tang , Shiyu Chang, Charu Aggarwal , Huan Liu. Negative Link Prediction in Social Media, Eighth ACM International Conference on Web Search and Data Mining . 02-FEB-15, . . . ,
- 11/17/2015 34.00 Xia Hu , Jiliang Tang, Yi Chang, Huan Liu. Predictability of Distrust with Interaction Data, 23rd ACM International Conference on Conference on Information and Knowledge Management. 03-NOV-14, . . . ,

TOTAL: 19

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

- | <u>Received</u> | <u>Paper</u> |
|------------------|--|
| 08/06/2013 13.00 | Pritam Gundecha, Geoffrey Barbier, Huan Liu. User Vulnerability and its Reduction on a Social Networking Site, Journals of Transactions on Knowledge Discovery from Data (08 2012) |
| 08/22/2012 4.00 | Pritam Gundecha, Huan Liu. Minimizing User Vulnerability and Retaining Social Utility in Social Media, Technical Report ASU CISE-2011-006 (11 2011) |
| 11/17/2015 24.00 | Pritam Gundecha, Huan Liu. Mining Social Media - A Brief Introduction, Tutorials in Operations Research (10 2012) |
| 11/17/2015 25.00 | Mohammad Ali Abbasi, Huan Liu, Reza Zafarani. Social Media Mining: Fundamental Issues and Challenges, IEEE International Conference on Data Mining (12 2013) |

TOTAL: 4

Number of Manuscripts:

Books

Received

Book

08/06/2013 12.00 Geoffrey Barbier, Zhuo Feng, Pritam Gundecha, Huan Liu. . Provenance Data in Social Media, San Rafael, California: Morgan & Claypool Publishers, (05 2013)

11/17/2015 29.00 Zhuo Feng, Pritam Gundecha, Huan Liu, Geoffrey Barbier. Provenance Data in Social Media, Synthesis Lectures on Data Mining and Knowledge Discovery: Morgan & Claypool Publishers, (05 2013)

11/17/2015 30.00 Reza Zafarani, Mohammad Ali Abbasi, Huan Liu. Social Media Mining: An Introduction, Cambridge University Press: Cambridge University Press, (05 2014)

11/17/2015 31.00 Huan Liu, Jiliang Tang. Trust in Social Media, Synthesis Lectures on Information Security, Privacy, and Trust: Morgan & Claypool Publishers, (09 2015)

TOTAL: 4

Received

Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

1. Jiliang Tang, ACM SIGKDD2015 Doctoral Dissertation Awards, Runner Up; Outstanding Computer Science PhD Student in CIDSE Award (May 1, 2015); Fulton Schools of Engineering Dean's Dissertation Award (May 13, 2015)
2. Xia Hu, Outstanding Graduate Student Award in Fulton Schools of Engineering, ASU (May 13, 2015); Outstanding Computer Science PhD Student in CIDSE Award (May 1, 2015)
3. Recipient of the "Best Researcher (Senior Faculty) Award (2014)", School of Computing, Informatics, and Decision Systems Engineering, ASU, May 1, 2015.
4. Recipient of the "Best Researcher (Senior Faculty) Award", School of Computing, Informatics, and Decision Systems Engineering, ASU, May 2, 2014.
5. The 2014 President's Award for Innovation at ASU, PI of Team Award on Empowering Humanitarian Assistance and Disaster Relief with Social Media and Data Analytics, April 2, 2014.
6. Recipient of the Yahoo Faculty Research and Engagement Program (FREPP), 2013
7. Recipient of the "Best Researcher (Senior Faculty) Award (2010)", Computer Science and Engineering, CIDSE, ASU, May 5, 2011.
8. IEEE Fellow, 2012
9. Recipient of the "Distinguished Contribution Award", PAKDD 2012, Kuala Lumpur, Malaysia, May 2012.
10. Mohammad-Ali Abbasi, Sun-Ki Chai, Huan Liu, and Kiran Sagoo. "Real-World Behavior Analysis through a Social Media Lens", International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction (SBP12), April 3-5, 2012. College Park, MD. Best Student Paper Award (Runner Up)

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Jiliang Tang	0.50	
Pritam Gundecha	0.50	
FTE Equivalent:	1.00	
Total Number:	2	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Huan Liu	0.10	
FTE Equivalent:	0.10	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Jiliang Tang

Pritam Gundecha

Total Number:

2

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Accomplishment 1: “Is Distrust the Negation of Trust? The Value of Distrust in Social Media”

- **Research Problem Studied:** Trust plays an important role in helping online users collect reliable information, and has attracted increasing attention in recent years. We learn from social sciences that, as the conceptual counterpart of trust, distrust could be as important as trust. However, little work exists in studying distrust in social media. What is the relationship between trust and distrust? Can we directly apply methodologies from social sciences to study distrust in social media? In this paper, we design two computational tasks by leveraging data mining and machine learning techniques to enable the computational understanding of distrust with social media data. The first task is to predict distrust from only trust, and the second task is to predict trust with distrust. We conduct experiments in real-world social media data. The empirical results of the first task provide concrete evidence to answer the question, “is distrust the negation of trust?” while the results of the second task help us figure out how valuable the use of distrust in trust prediction

- **Key Contributions:** As informed by social sciences, distrust could be as important as trust. A fundamental problem about distrust is what the relation between trust and distrust is. Passive observation is the modus operandi to obtain social media data, which lacks necessary information to apply methodologies from social sciences to understand distrust. However, an understanding of distrust with social media data is necessary because if distrust is the negation of trust, lacking distrust study matters little; while if distrust is a new dimension of trust, ignoring distrust in trust study may yield an incomplete and biased estimate of the effects of trust. In this paper, we first investigate the properties of distrust and find that we cannot equally and conversely extend the properties of trust to distrust. Then we then design two tasks by leveraging data mining and machine learning techniques to enable a computational understanding of distrust with social media data. The first task is to predict distrust with only trust information, and the second task is to predict trust with distrust information. We conduct experiments in real-world social media data. The evaluations of the first task suggests that distrust is not the negation of trust, while the results of the second task reveal that distrust has added value over trust.

Accomplishment 2: “Scalable Learning of Users’ Preferences Using Networked Data”

- **Research Problem Studied:** Users’ personal information such as their political views is important for many applications such as targeted advertisements or real-time monitoring of political opinions. Huge amounts of data generated by social media users present opportunities and challenge to study these preferences in a large scale. In this paper, we aim to infer social media users’ political views when only network information is available. In particular, given personal preferences about some of the social media users, how can we infer the preferences of unobserved individuals in the same network? There are many existing solutions that address the problem of classification with networked data problem. However, networks in social media normally involve millions and even hundreds of millions of nodes, which make the scalability an important problem in inferring personal preferences in social media. To address the scalability issue, we use social influence theory to construct new features based on a combination of local and global structures of the network. Then we use these features to train classifiers and predict users’ preferences. Due to the size of real-world social networks, using the entire network information is inefficient and not practical in many cases. By extracting local social dimensions, we present an efficient and scalable solution. Further, by capturing the network’s global pattern, the proposed solution, balances the performance requirement between accuracy and efficiency

- **Key Contributions:** In this paper we studied the network-based approach of inferring users’ personal preferences. We categorized the network-based algorithms into local and global algorithms. Local algorithms use users’ neighbors to predict their preferences, while the global approaches use the entire network information to predict user’s preferences. Our experimental results show that local algorithms are fast and scalable; however they need large amount of labeled data to achieve reasonable prediction accuracy. Further their prediction accuracy is always less than the accuracy of global algorithms. Global algorithms, in contrast, are computationally expensive, but perform well even in cases where only a very small fraction of the data is labeled. We proposed a new algorithm called LSocDim based on social influence theory to bridge the efficiency of local algorithms and the accuracy of global algorithms. The experiments show the efficiency and the effectiveness of the proposed algorithm. In particular, we show that LSocDim achieves a prediction accuracy near to that of the state-of-the-art global algorithm, SoCDim, while decreasing the running time by up to 40 times.

Accomplishment 3: “Leveraging Knowledge across Media for Spammer Detection in Microblogging”

- **Research Problem Studied:** While microblogging has emerged as an important information sharing and communication platform, it has also become a convenient venue for spammers to overwhelm other users with unwanted content. Currently, spammer detection in microblogging focuses on using social networking information, but little on content analysis due to the distinct nature of microblogging messages. First, label information is hard to obtain. Second, the texts in microblogging are short and noisy. As we know, spammer detection has been extensively studied for years in various media, e.g., emails, SMS and the web. Motivated by abundant resources available in the other media, we investigate whether we can take advantage of the existing resources for spammer detection in microblogging. While people accept that texts in microblogging are different from those in other media, there is no quantitative analysis to show how different they are. In this paper, we first perform a comprehensive linguistic study to compare spam across different media. Inspired by the findings, we present an optimization formulation that enables the design of spammer detection in microblogging using knowledge from external media. We conduct experiments on real-world Twitter datasets to verify (1) whether email, SMS and web spam resources help and (2) how different media help for spammer detection in microblogging.

- **Key Contributions:** Texts in microblogging are short, noisy, and labeling processing is time-consuming and labor-intensive, which presents great challenges for spammer detection. In this paper, we first conduct a quantitative analysis to study how noisy the microblogging texts are by comparing them with spam messages from other media. The results suggest that microblogging data is not significantly different from data from the other media. Based on the observations, a matrix factorization model is employed to learn lexicon information from external spam resources. By incorporating external

information from other media and content information from microblogging, we propose a novel framework for spammer detection. The experimental results demonstrate the effectiveness of our proposed model as well as the roles of different types of information in spammer detection.

Accomplishment 4: "Online Social Spammer Detection"

- **Research Problem Studied:** The explosive use of social media also makes it a popular platform for malicious users, known as social spammers, to overwhelm normal users with unwanted content. One effective way for social spammer detection is to build a classifier based on content and social network information. However, social spammers are sophisticated and adaptable to game the system with fast evolving content and network patterns. First, social spammers continually change their spamming content patterns to avoid being detected. Second, reflexive reciprocity makes it easier for social spammers to establish social influence and pretend to be normal users by quickly accumulating a large number of "human" friends. It is challenging for existing anti-spamming systems based on batch-mode learning to quickly respond to newly emerging patterns for effective social spammer detection. In this paper, we present a general optimization framework to collectively use content and network information for social spammer detection and provide the solution for efficient online processing. Experimental results on Twitter datasets confirm the effectiveness and efficiency of the proposed framework.

- **Key Contributions:** Social spammers are sophisticated and adaptable to game the system by continually change their content and network patterns. To handle fast evolving social spammers, we proposed to use online learning to efficiently reflect the newly emerging patterns. In this paper, we develop a general social spammer detection framework with both content and network information, and provide its online learning updating rules. In particular, we use directed graph Laplacian to model social network information, which is further integrated into a matrix factorization framework for content information modeling. By investigating its online updating scheme, we provide an efficient way for social spammer detection. Experimental results show that our proposed method is effective and efficient comparing with other social spammer detection methods.

Accomplishment 5: "Mining Social Media with Social Theories: A Survey"

- **Research Problem Studied:** The increasing popularity of social media encourages more and more users to participate in various online activities and produces data in an unprecedented rate. Social media data is big, linked, noisy, highly unstructured and incomplete, and differs from data in traditional data mining, which cultivates a new research field - social media mining. Social theories from social sciences are helpful to explain social phenomena. The scale and properties of social media data are very different from these of data social sciences use to develop social theories. As a new type of social data, social media data has a fundamental question - can we apply social theories to social media data? Recent advances in computer science provide necessary computational tools and techniques for us to verify social theories on large-scale social media data. Social theories have been applied to mining social media. In this article, we review some key social theories in mining social media, their verification approaches, interesting findings, and state-of-the-art algorithms. We also discuss some future directions in this active area of mining social media with social theories.

- **Key Contributions:** The social nature of social media data calls for new techniques and tools and cultivates a new field - social media mining. Social theories from social sciences have been proven to be applicable to mining social media. Integrating social theories with computational models is becoming an interesting way in mining social media data and makes exciting progress in various social media mining tasks. In this article, we review three key social theories, i.e., social correlation theory, balance theory and status theory, in mining social media data. In detail, we introduce basic concepts, verification methods, interesting findings and the state-of-the-art algorithms to exploit these social theories in social media mining tasks, which can be categorized to feature engineering, constraint generating and objective defining.

Accomplishment 6: "Social Recommendation: A Review"

- **Research Problem Studied:** Recommender systems play an important role in helping online users find relevant information by suggesting information of potential interest to them. Due to the potential value of social relations in recommender systems, social recommendation has attracted increasing attention in recent years. In this paper, we present a review of existing recommender systems and discuss some research directions. We begin by giving formal definitions of social recommendation and discuss the unique property of social recommendation and its implications compared with those of traditional recommender systems. Then, we classify existing social recommender systems into memory-based social recommender systems and model-based social recommender systems, according to the basic models adopted to build the systems, and review representative systems for each category. We also present some key findings from both positive and negative experiences in building social recommender systems, and research directions to improve social recommendation capabilities

- **Key Contributions:** Social recommendation has attracted broad attention from both academia and industry, and many social recommender systems have been proposed in recent years. In this paper, we first give a narrow definition and a broad definition of social recommendation to cover most existing definitions of social recommendation in literature, and discuss the unique feature of social recommender systems as well as its implications. We classify current social recommender systems into memory-based social recommender systems and model-based social recommender systems according to the basic models chosen to build the systems, and then present a review of representative systems for each category. We also discuss some key findings from positive and negative experiences in applying social recommender systems. Social recommendation is still in the early stages of development and needs further improvement. Finally we present research directions that can potentially improve performance of social recommender systems including exploiting the heterogeneity of social networks and weak dependence connections, microcosmic investigation of users and items, considering temporal information in rating and social information, understanding the role of negative relations, and integrating cross-media data.

Accomplishment 7: "Exploiting Homophily Effect for Trust Prediction"(hTrust)

- **Research Problem Studied:** Trust plays a crucial role for online users who seek reliable information. However, in reality,

user-specified trust relations are very sparse, i.e., a tiny number of pairs of users with trust relations are buried in a disproportionately large number of pairs without trust relations, making trust prediction a daunting task. As an important social concept, however, trust has received growing attention and interest. Social theories are developed for understanding trust. Homophily is one of the most important theories that explain why trust relations are established. Exploiting the homophily effect for trust prediction provides challenges and opportunities. In this paper, we embark on the challenges to investigate the trust prediction problem with the homophily effect. First, we delineate how it differs from existing approaches to trust prediction in an unsupervised setting. Next, we formulate the new trust prediction problem into an optimization problem integrated with homophily, empirically evaluate our approach on two datasets from real-world product review sites, and compare with representative algorithms to gain a deep understanding of the role of homophily in trust prediction.

- **Key Contributions:** In this paper, we study the problem of exploiting homophily effect for trust prediction. First we conduct experiments on datasets from real-world product review sites to demonstrate the existence of homophily in trust relations. Homophily regularization is then introduced to capture homophily effect in trust relations. An unsupervised framework is proposed, incorporating low-rank matrix factorization and homophily regularization. Extensive experiments are conducted to evaluate the proposed framework on real-world trust relation datasets and the experimental results demonstrate the effectiveness of our proposed framework as well as the role of homophily regularization for trust prediction.

Accomplishment 8: "Exploiting Local and Global Social Context for Recommendation"

- **Research Problem Studied:** With the fast development of social media, the information overload problem becomes increasingly severe and recommender systems play an important role in helping online users find relevant information by suggesting information of potential interests. Social activities for online users produce abundant social relations. Social relations provide an independent source for recommendation, presenting both opportunities and challenges for traditional recommender systems. Users are likely to seek suggestions from both their local friends and users with high global reputations, motivating us to exploit social relations from local and global perspectives for online recommender systems in this paper. We develop approaches to capture local and global social relations, and propose a novel framework LOCABAL taking advantage of both local and global social context for recommendation. Empirical results on real-world datasets demonstrate the effectiveness of our proposed framework and further experiments are conducted to understand how local and global social context work for the proposed framework.

- **Key Contributions:** The availability of social relations presents both challenges and opportunities for traditional recommender systems. In this paper, we investigate how to exploit local and global social context for recommendation. To capture local social context, we force that the user preferences of two socially connected users are correlated as suggested by social correlation theories and we also study the connections between our proposed approach and existing approaches. Ratings from users with high reputations are more likely to be trustworthy; therefore, to capture global social context, we use user reputation scores to weight the importance of their ratings. With these solutions, we propose a framework LOCABAL to integrate local and global social context for recommendation. Experimental results on real-world data sets show that the proposed framework LOCABAL outperforms representative social recommender systems. Further experiments are conducted to understand the working of LOCABAL.

Accomplishment 9: "A Tool for Collecting Provenance Data in Social Media"

- **Research Problem Studied:** In recent years, social media sites have provided a large amount of information. Recipients of such information need mechanisms to know more about the received information, including the provenance. Previous research has shown that some attributes related to the received information provide additional context, so that recipient can assess the amount of value, trust, and validity to be placed in the received information. Personal attributes of a user, including name, location, education, ethnicity, gender, and political and religious affiliations, can be found in social media sites. In this paper, we present a novel web-based tool for collecting the attributes of interest associated with a particular social media user related to the received information. This tool provides a way to combine different attributes available at different social media sites into a single user profile. Using different types of Twitter users, we also evaluate the performance of the tool in terms of number of attribute values collected, validity of these values, and total amount of retrieval time.

- **Key Contributions:** The provenance data collector tool aims to collect provenance attribute values of a user. By collecting such values of a user related to the received information, the tool could facilitate recipients to understand more about the received information. Data generated on social media sites is largely distributed and unstructured in nature. The proposed tool also provides a way to combine such distributed and unstructured social media data.

Accomplishment 10: "Recovering Information Recipients in Social Media via Provenance"

- **Research Problem Studied:** In recent years, social media has changed the way we interact and communicate. Although the existing structure of social media allows users to easily create, receive, and propagate pieces of information, many a time, users do not have background knowledge about the received information, including the provenance (sources or originators) of information, and other recipients who may have retransmitted or modified the information. Providing such additional context to the received information can help users know how much value, trust, and validity should be placed in received information. To judge the credibility of the received piece of information, it is vital to know who are its sources, and how information propagates from sources to other social media users. In this paper, we are studying a novel research problem that facilitates a few known recipients to recover other unknown recipients, and seek the provenance of information. The experimental results with Facebook and Twitter datasets show that the proposed algorithm is effective in correctly recovering the unknown recipients and seeking the provenance of information.

- **Key Contributions:** Social media allows its users to share a vast amount of information with other users, but it provides no mechanism to know more about the received information for its users. In this paper, we aim to recover information recipients as

well as seek the provenance by knowing a few nodes and using only link information in social networks. Information recipients exist along the paths from the sources to the known nodes. In this paper we seek the information propagation flow from the sources to the known nodes, and recover the most likely information recipients. Using the experiment results from the Facebook and Twitter datasets, we show that the proposed algorithm is effective in correctly recovering the information recipients and seeking the provenance of information.

Accomplishment 11: "Context-Aware Review Helpfulness Rating Prediction"

- **Research Problem Studied:** Online reviews play a vital role in the decision-making process for online users. Helpful reviews are usually buried in a large number of unhelpful reviews, and with the consistently increasing number of reviews, it becomes more and more difficult for online users to find helpful reviews. Therefore most online review websites allow online users to rate the helpfulness of a review and a global helpfulness score is computed for the review based on its available ratings. However, in reality, user-specified helpfulness ratings for reviews are very sparse a few reviews attract large numbers of helpfulness ratings while most reviews obtain few or even no helpfulness ratings. The available helpfulness ratings are too sparse for online users to assess the helpfulness of reviews. Also the helpfulness of a review is not necessarily equally useful for all users and users with different background may treat the helpfulness of a review very differently. The user idiosyncrasy of review helpfulness motivates us to study the problem of review helpfulness rating prediction in this paper. We first identify various types of context information, model them mathematically, and propose a context-aware review helpfulness rating prediction framework CAP. Experimental results demonstrate the effectiveness of the proposed framework and the importance of context awareness in solving the review helpfulness rating prediction problem.

- **Key Contributions:** In this paper we study the problem of review helpfulness rating prediction by exploiting context awareness to infer unknown helpfulness ratings automatically, motivated by the fact that helpful reviews can be buried in large amounts of useless reviews and the user-specific helpfulness ratings are too sparse for online users to assess the helpfulness of reviews. We first show that the problem we study differs from review quality prediction problem and the item rating prediction problem. We extract four types of social context, i.e., author context, rater context, connection context and preference context, formulate them mathematically, and propose a context-aware helpfulness prediction framework CAP which exploits content context and various types of social context. Experimental results demonstrate that our proposed framework outperforms the state-of-the-art baseline methods with both cold-start and warm-start settings, and further experiments are conducted to understand the importance of context awareness in the proposed framework.

Accomplishment 12: "Seeking Provenance of Information in Social Media"

- **Research Problem Studied:** Social media has profoundly impacted the way people interact and communicate. Social media propagates breaking news and disinformation alike fast and on an unsurpassed scale. Because of its democratization nature, social media users can easily produce, receive and propagate a piece of information without necessarily providing traceable information. Thus, there are no means for a user to verify the provenance (also known as, sources or originators) of information. The disinformation can cause tragic consequences to society and individuals. This work aims to take advantage of characteristics of social media to provide a solution to the problem of lacking traceable information. Such knowledge can provide additional context to the received information such that a user can assess how much value, trust, and validity should be placed in received information. In this paper, we are studying a novel research problem that facilitates a few known recipients (less than 1% of the total recipients) to seek the provenance of information by recovering how it has own from its originators. The proposed methodology exploits easily computable node centralities of a large social media network. The experimental results with Facebook and Twitter datasets show that the proposed mechanism is effective in correctly identifying the additional recipients and seeking the provenance of information.

- **Key Contributions:** Social media allows its users to share vast amount of information with other users, but it lacks mechanisms that provide traceable knowledge about the received information for its users. In this paper, we study a novel research problem that facilitates a few P-nodes (less than 1% of total recipients) to seek the provenance of information by identifying how it has own from its originators. To this end, we first formally present the problem and provide the complexity analysis. Then, use the Facebook and Twitter datasets to show the existence of two hypotheses: Degree Propensity and Closeness Propensity. The proposed methodology then exploits these hypotheses to provide not only the critical information about the provenance, but also the most likely provenance paths. Finally using the experimental results with the Facebook and Twitter datasets, we show that the proposed algorithm is effective in correctly identifying the additional transmitters, and seeking the provenance of information.

Accomplishment 13: "A Tool for Assisting Provenance Search in Social Media"

- **Research Problem Studied:** In recent years, social media sites are witnessing an information explosion. Determining the reliability of such a large amount of information is a major area of research. Information provenance (aka, sources or origin) provides a way to measure the reliability of information in social networks. The main challenge in seeking provenance is the availability of suitable data consisting of sufficient unique propagation paths. Current research on provenance in social media uses synthetically generated propagation paths. Although these proposed approaches are theoretically significant, it is still a challenge to apply and evaluate them on social media. Hence, knowledge of the actual propagation paths for a piece of information will be a valuable asset in provenance search. This paper presents a tool for capturing the propagation network of a given tweet or URL (Uniform Resource Locator) in the Twitter network. Researchers can use this tool to collect information propagation data, design effective strategies for determining the provenance, and gain information about the tweet such as impact, growth rate and users influencing the spread. An overview of the user interface and the architecture of the system is provided. Two case studies, one relating to disinformation in riot situations and another on corporate involvement in education has been presented to demonstrate the effectiveness of the system for seeking provenance information.

- **Key Contributions:** The paper presents a tool to obtain the spread of a given tweet or URL on the twitter network. The tool presents researchers with a propagation network to assist in seeking the provenance path of a given tweet. The provenance path gives additional information to assess the reliability of a given piece of data in social media.

Accomplishment 14: "Provenance Data in Social Media"

- **Book Overview:** Social media shatters the barrier to communicate anytime anywhere for people of all walks of life. The publicly available, virtually free information in social media poses a new challenge to consumers who have to discern whether a piece of information published in social media is reliable. For example, it can be difficult to understand the motivations behind a statement passed from one user to another, without knowing the person who originated the message. Additionally, false information can be propagated through social media, resulting in embarrassment or irreversible damages. Provenance data associated with a social media statement can help dispel rumors, clarify opinions, and confirm facts. However, provenance data about social media statements is not readily available to users today. Currently, providing this data to users requires changing the social media infrastructure or offering subscription services. Taking advantage of social media features, research in this nascent field spearheads the search for a way to provide provenance data to social media users, thus leveraging social media itself by mining it for the provenance data. Searching for provenance data reveals an interesting problem space requiring the development and application of new metrics in order to provide meaningful provenance data to social media users. This lecture reviews the current research on information provenance, explores exciting research opportunities to address pressing needs, and shows how data mining can enable a social media user to make informed judgements about statements published in social media.

- **Table of Contents:** Information Provenance in Social Media / Provenance Attributes / Provenance via Network Information / Provenance Data

Accomplishment 15: "User Vulnerability and its Reduction on a Social Networking Site"

- **Research Problem Studied:** Privacy and security are major concerns for many users of social media. When users share information (e.g., data and photos) with friends, they can make their friends vulnerable to security and privacy breaches with dire consequences. With the continuous expansion of a user's social network, privacy settings alone are often inadequate to protect user's profile. In this research, we aim to address some critical issues related to privacy protection: (1) How can we measure and assess individual user's vulnerability? (2) With the diversity of one's social network friends, how can one figure out an effective approach to maintaining balance between vulnerability and social utility? In this work, first we present a novel way to define vulnerable friends from an individual user's perspective. User vulnerability is dependent on whether or not the user's friends' privacy settings protect the friend and the individual's network of friends (which includes the user). We show that it is feasible to measure and assess user vulnerability, and reduce one's vulnerability without changing the structure of a social networking site. The approach is to unfriend one's most vulnerable friends. However, when such a vulnerable friend is also socially important, unfriending him would significantly reduce one's own social status. We formulate this novel problem as vulnerability minimization with social utility constraints. We formally define the optimization problem, and provide an approximation algorithm with a proven bound. Finally, we conduct a large-scale evaluation of new framework using a Facebook dataset. We resort to experiments and observe how much vulnerability an individual user can decrease by unfriending a vulnerable friend. We compare performance of different unfriending strategies and discuss the security risk of new friend request. Additionally, by employing different forms of social utility, we confirm that balance between user vulnerability and social utility can be practically achieved.

- **Key Contributions:** We propose a feasible approach to a novel problem of identifying a user's vulnerable friends on a social networking site. Our work differs from existing work addressing social networking privacy by introducing a vulnerability-centered approach to a user security on a social networking site. On most social networking sites, privacy related efforts have been concentrated on protecting individual attributes only. However, users are often vulnerable through community attributes. Unfriending vulnerable friends can help protect users against the security risks. Based on our study of over 2 million users, we find that users are either not careful or not aware of security and privacy concerns of their friends. Our model clearly highlights the impact of each new friend on a user's privacy. Our approach does not require the structural change of a social networking site and aims to maximally reduce a user's vulnerability while minimizing his social utility loss. The work formulates a novel problem of constrained vulnerability reduction suggests a feasible approach, and demonstrates that the problem of constrained vulnerability reduction is solvable.

Accomplishment 16: "mTrust: Discerning Multi-Faceted Trust in a Connected World"

- **Research Problem Studied:** The issue of trust has attracted increasing attention from the community of social media research. Trust, as a social concept, naturally has multiple facets, indicating multiple and heterogeneous trust relationships between users. Here is a multifaceted trust example from Epinions. Figure 1(a) shows single trust relationships between user 1 and his 20 friends. Here, we can see that user 7 is the more trustable for user 1. Figures 1(b) and 1(c) show their multifaceted trust relationships in the categories "home and garden" and "restaurants" respectively. For the category "home and garden" user 7 is not necessary the most trusted friend of user 1. This shows that trust relationships in different categories vary. Thus, people trust others differently in different facets.

(a) Single Trust (b) Trust in home and garden (c) Trust in restaurants

Figure 1: Single trust and multifaceted trust relationships of one use in Epinions.

(Note: The thickness of a line indicates the level of trust.)

There are two challenges to study in obtaining multifaceted trust between users: first, the representation of multiple and

heterogeneous trust relationships between users, and second, estimating the strength of multifaceted trust. Traditionally, trust is represented by an adjacency matrix. However, this cannot capture the multifaceted trust relations. We developed a new algorithm, mTrust, which extends a matrix representation to a tensor representation, adding an extra dimension for facet description. Previous work observed a strong correlation between trust and user similarity in the context of rating systems. Therefore, it is reasonable to embed trust strength inference in rating prediction. Thus, to evaluate the usefulness of multifaceted trust, this work embeds the multifaceted trust inference in the framework of rating prediction.

- **Key Contributions:** Interesting findings from the experiments are that (1) more than 20% of reciprocal links are heterogeneous, (2) more than 14% transitive trust relations are heterogeneous, and (3) more than 11% of cocitation trust relations are heterogeneous. With these findings, mTrust can be applied to many online tasks such as improving rating prediction, enabling facet-sensitive ranking, and making status theory applicable to reciprocal links.

Accomplishment 17: “eTrust: Understanding Trust Evolution in an Online World”

- **Research Problem Studied:** Most existing research about online trust assumes static trust relations between users. As we are informed by social sciences, trust evolves as humans interact. Little work exists studying trust evolution in an online world. Researching online trust evolution faces unique challenges because more often than not, available data is from passive observation. In this paper, we leverage social science theories to develop a methodology that enables the study of online trust evolution. In particular, we propose a framework of evolution trust, eTrust, which exploits the dynamics of user preferences in the context of online product review. We present technical details about modeling trust evolution, and perform experiments to show how the exploitation of trust evolution can help improve the performance of online applications such as rating and trust prediction.

- **Key Contributions:** We study online trust evolution in the context of product review. By exploiting the correlation between user preferences and trust relations, we propose a framework, eTrust, to understand the evolution of trust in an online world and apply eTrust to various online applications such as rating prediction and trust prediction. Interesting findings are observed in our experiments using real-world data, Epinions; and eTrust can be applied to improve the performance of rating prediction and trust prediction.

Accomplishment 18: “Minimizing User Vulnerability and Retaining Social Utility in Social Media”

- **Research Problem Studied:** Privacy and security are major concerns for many users of social media. When users share information (e.g., data and photos) with friends, they can make their friends vulnerable to security and privacy breach with dire consequences. In our earlier work, we show that it is feasible to measure user vulnerability and reduce one's vulnerability without changing the structure of a social networking site. The approach is to unfriend one's most vulnerable friends. However, when such a vulnerable friend is also socially important, unfrinding him would significantly reduce one's own social status. In this work, we address the problem of vulnerability minimization with minimum social utility losses. This work extends the existing vulnerability reduction model to a more general form. Using a general model, we formulate the two discrete optimization problems. Both problems are NP-hard.

- **Key Contributions:** We formally formulate the optimization problem, propose an approximation algorithm with a proven bound, and conduct empirical experiments with different forms of social utility on a large-scale Facebook dataset for performance evaluation and comparison. Our work differs from existing work addressing social networking privacy. Our approach does not require the structural change of a social networking site and aims to maximally reduce a user's vulnerability while minimizing his social utility loss.

Technology Transfer

Final Report: Assessing Trustworthiness in Social Media: A Social Computing Approach, ARO (#025071)

Assessing Trustworthiness in Social Media A Social Computing Approach

Project Description

Social media is gaining popularity in recent years and increasingly becoming an integral part of our life. Given the extensiveness, instantaneity, and diffusion speed of social media, e.g., a tweet or a clip of video, could galvanize a digital revolution or wreak havoc with one's otherwise routine and uneventful working life. With the presence of adversaries, the convenient use of and low barrier of social media brings about new challenges. How well we address these challenges can directly influence our ability to manage information and misinformation, and the future role of social media as a reliable communication mechanism. One such pressing challenge is to *assess information trustworthiness in social media*. We propose to investigate research issues related to social media trustworthiness and its assessment by leveraging social research methods, developing new computational social methods, and creating novel approaches to social media data collection and sharing.

Research Problem

In social sciences, trust is about a relationship between two entities, the trustor and the trustee. Trust can be defined as the perception of the trustor about the degree to which the trustee would satisfy an expectation. Trustworthiness can be defined from the perspective of both entities; in this work, it is the perspective of the trustor that defines a property that can be judged, i.e., the amount of trust associated with the trustee. In all cases trust is a heuristic decision rule, allowing the human to deal with complexities that would require unrealistic effort in rational reasoning. One of the key current challenge is to rethink how the rapid progress of technology has impacted trust as information technology has significantly changed how people interact, express themselves, and behave. The assessment of information trustworthiness in social media requires answers to the three essential questions about the information: (1) source (or author), (2) author position, and (3) content. The search for the answers is greatly complicated by the nature of social media: enormous sizes in terms of users and links, irregular uses of languages, incomplete sentences or messages, and inordinate amounts of data and meta-data. In addition, both linked data and attribute data are present in social media. The former represents the connectedness among entities and the latter the properties of entities. In search of the three answers, we face research challenges:

1. Information Provenance - Identifying the true source (or author) of information,
2. Friendship Differentiation - Determining if the author is a friend, acquaintance, or foe, and
3. Content Analysis - Analyzing the content to ascertain its intention, quality, and etc.

In this project, we focus on developing computational social theories and methods for the first two challenges. The third challenge is partially addressed in our recent work. Additional work on trust maintenance can be found in literature.

Subject Terms

Social Networks, Social Media, Social Media Mining, Trust, Information Provenance.

Publications

- **PhD Dissertations**
 - Pritam Gundecha. "Manging a Uer's Vulneratilly on a Social Networking Site", Computer Science & Engineering, Arizona State University, March, 2015.
 - Jiliang Tan. "Computing Distrust in Social Media", Computer Science & Engineering, Arizona State University, February, 2015.
- **Books**
 - Jiliang Tang and Huan Liu. "Trust in Social Media", Morgan & Claypool Publishers, August, 2015
 - Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. "Social Media Mining: An Introduction", Cambridge University Press, May, 2014.
 - Geoffrey Barbier, Zhuo Feng, Pritam Gundecha, and Huan Liu. "Provenance Data in Social Media", Synthesis Lectures on Data Mining and Knowledge Discovery, Morgan & Claypool Publishers, 2013.
- **Tutorials**
 - Dongwon Lee and Huan Liu. "LIKE and Recommendation in Social Media", International Conference on World Wide Web (WWW), April, 2015.
 - Jiliang Tang, Jie Tang and Huan Liu. "Recommendation in Social Media - Recent Advances and New Frontier", 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), August, 2014.
 - Jiliang Tang and Huan Liu. "Trust in Social Computing", International Conference on World Wide Web (WWW), April, 2014.
 - Mohammad Ali Abbasi, Huan Liu, and Reza Zafarani. "Social Media Mining: Fundamental Issues and Challenges", ICDM2013.
 - Pritam Gundecha and Huan Liu "Mining Social Media - A Brief Introduction", INFORMS 2012.
- **Journal Articles**
 - Jiliang Tang, Huiji Gao, Atish Dasarma, Yingzhou Bi, and Huan Liu. "Trust Evolution: Modeling and Its Applications", IEEE Transactions on Knowledge and Data Engineering (TKDE).27(6): 1724-1738 (2015).

- Pritam Gundecha, Geoffrey Barbier, Jiliang Tang and Huan Liu. "User Vulnerability and its Reduction on a Social Networking Site", *Journals of Transactions on Knowledge Discovery from Data*, (TKDD 2014), 9(2): Pages 12:1-12:25, 2014.
- Jiliang Tang, Yi Chang and Huan Liu. "Mining Social Media with Social Theories: A Survey", *SIGKDD Explorations*, 15(2): 20 - 29. December 2013.
- Jiliang Tang, Xia Hu and Huan Liu. "Social Recommendation: A Review", *Social Network Analysis and Mining*, Springer. 3(4): 1113-1133 (2013).
- **Conferences and Workshops**
 - Jiliang Tang, Shiyu Chang, Charu Aggarwal, and Huan Liu. "Negative Link Prediction in Social Media". *ACM International Conference on Web Search and Data Mining (WSDM2015)*, February 2-6, 2015. Shanghai, China.
 - Jiliang Tang, Xia Hu, Yi Chang, and Huan Liu. "Predictability of Distrust with Interaction Data", *ACM International Conference of Information and Knowledge Management (CIKM2014)*, November 3-7, 2014. Shanghai, China.
 - Jiliang Tang, Xia Hu and Huan Liu. "Is Distrust the Negation of Trust? The Value of Distrust in Social Media", *ACM Hypertext conference*, 2014.
 - Mohammad Ali Abbasi, Jiliang Tang, and Huan Liu. "Scalable Learning of Users' Preferences Using Networked Data". *ACM Hypertext conference*, 2014.
 - Xia Hu, Jiliang Tang, and Huan Liu. "Leveraging Knowledge across Media for Spammer Detection in Microblogging", the 37th Annual ACM Special Interest Group on Information Retrieval Conference (SIGIR 2014).
 - Xia Hu, Jiliang Tang, and Huan Liu. "Online Social Spammer Detection". the 28th AAAI Conference on Artificial Intelligence (AAAI 2014).
 - Suhas Ranganath, Pritam Gundecha, and Huan Liu. "A Tool for Assisting Provenance Search in Social Media", demonstration paper, the 22nd ACM International Conference on Information and Knowledge Management (CIKM 2013). (To Appear)
 - Pritam Gundecha, Zhuo Feng, and Huan Liu. "Seeking Provenance of Information in Social Media", short paper, the 22nd ACM International Conference on Information and Knowledge Management (CIKM 2013). (To Appear)
 - Jiliang Tang, Huiji Gao, Xia Hu, and Huan Liu. "Context-Aware Review Helpfulness Rating Prediction", the 7th ACM Recommender Systems Conference (RecSys 2013). (To Appear)
 - Zhuo Feng, Pritam Gundecha, and Huan Liu. "Recovering Information Recipients in Social Media via Provenance", short paper, the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013). [\[pdf\]](#)
 - Pritam Gundecha, Suhas Ranganath, Zhuo Feng, and Huan Liu. "A Tool for Collecting Provenance Data in Social Media", demonstration paper, the 19th ACM SIGKDD International Conference On Knowledge Discovery and Data Mining (SIGKDD 2013). [\[pdf\]](#) [\[video\]](#) [\[Provenance Data Collector\]](#)
 - Jiliang Tang, Xia Hu, Huiji Gao, and Huan Liu. "Exploiting Local and Global Social Context for Recommendation", the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013). [\[pdf\]](#)

- Jiliang Tang, Huiji Gao, Xia Hu, and Huan Liu. "Exploiting Homophily Effect for Trust Prediction"(hTrust), the 6th ACM International Conference on Web Search and Data Mining (WSDM 2013). [\[pdf\]](#) [\[PPT\]](#) [\[Poster\]](#) {DOI:10.1145/2433396.2433405}
- Jiliang Tang, Huiji Gao, Huan Liu and Atish Das Sarma. "eTrust: Understanding Trust Evolution in an Online World", the Eighteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining , 2012. [\[pdf\]](#) [\[video\]](#) {DOI: 10.1145/2339530.2339574}
- Jiliang Tang, Huiji Gao, and Huan Liu. "mTrust: Discerning Multi-Faceted Trust in a Connected World", the 5th ACM International Conference on Web Search and Data Mining (WSDM 2012), February 8-12, 2012. Seattle, Washington [\[pdf\]](#) [\[PPT\]](#) {DOI: 10.1145/2124295.2124309}
- **Books or Chapters**
 - Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. "Social Media Mining: An Introduction", Cambridge University Press, ISBN: 9781107018853. May, 2014
 - Geoffrey Barbier, Zhuo Feng, Pritam Gundecha, and Huan Liu. "Provenance Data in Social Media", Synthesis Lectures on Data Mining and Knowledge Discovery, Morgan & Claypool Publishers, 2013. [\[link\]](#)
 - Pritam Gundecha, and Huan Liu. "Mining Social Media : A Brief Introduction", the Institute for Operations Research and the Management Sciences (INFORMS) TutORials in Operations Research, 2012. [\[pdf\]](#) {DOI: 10.1287/educ.1120.0105}
- **Technical Reports**
 - Pritam Gundecha, and Huan Liu. "Minimizing User Vulnerability and Retaining Social Utility in Social Media", ASUCISE-2011-006, School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, AZ 85287. Nov. 2011. [\[pdf\]](#)

Honors and Awards

- Recipient of the Yahoo Faculty Research and Engagement Program (FREP), 2013
- Recipient of the "Best Researcher (Senior Faculty) Award (2010)", Computer Science and Engineering, CIDSE, ASU, May 5, 2011.
- IEEE Fellow, 2012
- Recipient of the "Distinguished Contribution Award", PAKDD 2012, Kuala Lumpur, Malaysia, May 2012.
- Mohammad-Ali Abbasi, Sun-Ki Chai, Huan Liu, and Kiran Sagoo. "Real-World Behavior Analysis through a Social Media Lens", International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction (SBP12), April 3-5, 2012. College Park, MD. Best Student Paper Award (Runner Up)

Related News

- Invited to join live panel discussion on “Selling Personal Data” on HuffPost Live, Huffington Post, July 30.
- Scholarship to attend the "Summer School on Formal Methods for the Science of Security", July 22-26 at the UIUC.

Resources and Project Website

- Provenance Paths Collector (CIKM 2013) [[Tool](#)] [[Demo](#)]
- Provenance Attributes Collector (KDD 2013) [[Tool](#)] [[Demo](#)]
- [Product Review Datasets\(Epinions and Ciao\)](#)
- Project website: <http://www.public.asu.edu/~huanliu/projects/ARO11/>

Project Members (current and former)

- [Huan Liu](#) (PI)
- [Pritam Gundecha](#) (PhD Student)
- [Jiliang Tang](#)(PhD Student)
- Soumya Kambhampati (High School Student)
- Terry Wen (High School Student)

Appendix for Scientific Progress

Accomplishment 1: “Is Distrust the Negation of Trust? The Value of Distrust in Social Media”

- Research Problem Studied: Trust plays an important role in helping online users collect reliable information, and has attracted increasing attention in recent years. We learn from social sciences that, as the conceptual counterpart of trust, distrust could be as important as trust. However, little work exists in studying distrust in social media. What is the relationship between trust and distrust? Can we directly apply methodologies from social sciences to study distrust in social media? In this paper, we design two computational tasks by leveraging data mining and machine learning techniques to enable the computational understanding of distrust with social media data. The first task is to predict distrust from only trust, and the second task is to predict trust with distrust. We conduct experiments in real-world social media data. The empirical results of the first task provide concrete evidence to answer the question, “is distrust the negation of trust?” while the results of the second task help us figure out how valuable the use of distrust in trust prediction
- Key Contributions: As informed by social sciences, distrust could be as important as trust. A fundamental problem about distrust is what the relation between trust and distrust is. Passive observation is the modus operandi to obtain social media data, which lacks necessary information to apply methodologies from social sciences to understand distrust. However, an understanding of distrust with social media data is necessary because if

distrust is the negation of trust, lacking distrust study matters little; while if distrust is a new dimension of trust, ignoring distrust in trust study may yield an incomplete and biased estimate of the effects of trust. In this paper, we first investigate the properties of distrust and find that we cannot equally and conversely extend the properties of trust to distrust. Then we then design two tasks by leveraging data mining and machine learning techniques to enable a computational understanding of distrust with social media data. The first task is to predict distrust with only trust information, and the second task is to predict trust with distrust information. We conduct experiments in real-world social media data. The evaluations of the first task suggests that distrust is not the negation of trust, while the results of the second task reveal that distrust has added value over trust.

Accomplishment 2: “Scalable Learning of Users' Preferences Using Networked Data”

- **Research Problem Studied:** Users' personal information such as their political views is important for many applications such as targeted advertisements or real-time monitoring of political opinions. Huge amounts of data generated by social media users present opportunities and challenge to study these preferences in a large scale. In this paper, we aim to infer social media users' political views when only network information is available. In particular, given personal preferences about some of the social media users, how can we infer the preferences of unobserved individuals in the same network? There are many existing solutions that address the problem of classification with networked data problem. However, networks in social media normally involve millions and even hundreds of millions of nodes, which make the scalability an important problem in inferring personal preferences in social media. To address the scalability issue, we use social influence theory to construct new features based on a combination of local and global structures of the network. Then we use these features to train classifiers and predict users' preferences. Due to the size of real-world social networks, using the entire network information is inefficient and not practical in many cases. By extracting local social dimensions, we present an efficient and scalable solution. Further, by capturing the network's global pattern, the proposed solution, balances the performance requirement between accuracy and efficiency
- **Key Contributions:** In this paper we studied the network-based approach of inferring users' personal preferences. We categorized the network-based algorithms into local and global algorithms. Local algorithms use users' neighbors to predict their preferences, while the global approaches use the entire network information to predict user's preferences. Our experimental results show that local algorithms are fast and scalable; however they need large amount of labeled data to achieve reasonable prediction accuracy. Further their prediction accuracy is always less than the accuracy of global algorithms. Global algorithms, in contrast, are computationally expensive, but perform well even in cases where only a very small fraction of the data is labeled. We proposed a new algorithm called LSocDim based on social influence theory to bridge the efficiency

of local algorithms and the accuracy of global algorithms. The experiments show the efficiency and the effectiveness of the proposed algorithm. In particular, we show that LSocDim achieves a prediction accuracy near to that of the state-of-the-art global algorithm, SoCDim, while decreasing the running time by up to 40 times.

Accomplishment 3: “Leveraging Knowledge across Media for Spammer Detection in Microblogging”

- **Research Problem Studied:** While microblogging has emerged as an important information sharing and communication platform, it has also become a convenient venue for spammers to overwhelm other users with unwanted content. Currently, spammer detection in microblogging focuses on using social networking information, but little on content analysis due to the distinct nature of microblogging messages. First, label information is hard to obtain. Second, the texts in microblogging are short and noisy. As we know, spammer detection has been extensively studied for years in various media, e.g., emails, SMS and the web. Motivated by abundant resources available in the other media, we investigate whether we can take advantage of the existing resources for spammer detection in microblogging. While people accept that texts in microblogging are different from those in other media, there is no quantitative analysis to show how different they are. In this paper, we first perform a comprehensive linguistic study to compare spam across different media. Inspired by the findings, we present an optimization formulation that enables the design of spammer detection in microblogging using knowledge from external media. We conduct experiments on real-world Twitter datasets to verify (1) whether email, SMS and web spam resources help and (2) how different media help for spammer detection in microblogging.
- **Key Contributions:** Texts in microblogging are short, noisy, and labeling processing is time-consuming and labor-intensive, which presents great challenges for spammer detection. In this paper, we first conduct a quantitative analysis to study how noisy the microblogging texts are by comparing them with spam messages from other media. The results suggest that microblogging data is not significantly different from data from the other media. Based on the observations, a matrix factorization model is employed to learn lexicon information from external spam resources. By incorporating external information from other media and content information from microblogging, we propose a novel framework for spammer detection. The experimental results demonstrate the effectiveness of our proposed model as well as the roles of different types of information in spammer detection.

Accomplishment 4: “Online Social Spammer Detection”

- **Research Problem Studied:** The explosive use of social media also makes it a popular platform for malicious users, known as social spammers, to overwhelm normal users with unwanted content. One effective way for social spammer detection is to build a classifier

based on content and social network information. However, social spammers are sophisticated and adaptable to game the system with fast evolving content and network patterns. First, social spammers continually change their spamming content patterns to avoid being detected. Second, reflexive reciprocity makes it easier for social spammers to establish social influence and pretend to be normal users by quickly accumulating a large number of “human” friends. It is challenging for existing anti-spamming systems based on batch-mode learning to quickly respond to newly emerging patterns for effective social spammer detection. In this paper, we present a general optimization framework to collectively use content and network information for social spammer detection and provide the solution for efficient online processing. Experimental results on Twitter datasets confirm the effectiveness and efficiency of the proposed framework.

- **Key Contributions:** Social spammers are sophisticated and adaptable to game the system by continually change their content and network patterns. To handle fast evolving social spammers, we proposed to use online learning to efficiently reflect the newly emerging patterns. In this paper, we develop a general social spammer detection framework with both content and network information, and provide its online learning updating rules. In particular, we use directed graph Laplacian to model social network information, which is further integrated into a matrix factorization framework for content information modeling. By investigating its online updating scheme, we provide an efficient way for social spammer detection. Experimental results show that our proposed method is effective and efficient comparing with other social spammer detection methods.

Accomplishment 5: “Mining Social Media with Social Theories: A Survey”

- **Research Problem Studied:** The increasing popularity of social media encourages more and more users to participate in various online activities and produces data in an unprecedented rate. Social media data is big, linked, noisy, highly unstructured and incomplete, and differs from data in traditional data mining, which cultivates a new research field - social media mining. Social theories from social sciences are helpful to explain social phenomena. The scale and properties of social media data are very different from these of data social sciences use to develop social theories. As a new type of social data, social media data has a fundamental question - can we apply social theories to social media data? Recent advances in computer science provide necessary computational tools and techniques for us to verify social theories on large-scale social media data. Social theories have been applied to mining social media. In this article, we review some key social theories in mining social media, their verification approaches, interesting findings, and state-of-the-art algorithms. We also discuss some future directions in this active area of mining social media with social theories.
- **Key Contributions:** The social nature of social media data calls for new techniques and tools and cultivates a new field - social media mining. Social theories from social sciences have been proven to be applicable to mining social media. Integrating social

theories with computational models is becoming an interesting way in mining social media data and makes exciting progress in various social media mining tasks. In this article, we review three key social theories, i.e., social correlation theory, balance theory and status theory, in mining social media data. In detail, we introduce basic concepts, verification methods, interesting findings and the state-of-the-art algorithms to exploit these social theories in social media mining tasks, which can be categorized to feature engineering, constraint generating and objective defining.

Accomplishment 6: "Social Recommendation: A Review"

- **Research Problem Studied:** Recommender systems play an important role in helping online users find relevant information by suggesting information of potential interest to them. Due to the potential value of social relations in recommender systems, social recommendation has attracted increasing attention in recent years. In this paper, we present a review of existing recommender systems and discuss some research directions. We begin by giving formal definitions of social recommendation and discuss the unique property of social recommendation and its implications compared with those of traditional recommender systems. Then, we classify existing social recommender systems into memory-based social recommender systems and model-based social recommender systems, according to the basic models adopted to build the systems, and review representative systems for each category. We also present some key findings from both positive and negative experiences in building social recommender systems, and research directions to improve social recommendation capabilities
- **Key Contributions:** Social recommendation has attracted broad attention from both academia and industry, and many social recommender systems have been proposed in recent years. In this paper, we first give a narrow definition and a broad definition of social recommendation to cover most existing definitions of social recommendation in literature, and discuss the unique feature of social recommender systems as well as its implications. We classify current social recommender systems into memory-based social recommender systems and model-based social recommender systems according to the basic models chosen to build the systems, and then present a review of representative systems for each category. We also discuss some key findings from positive and negative experiences in applying social recommender systems. Social recommendation is still in the early stages of development and needs further improvement. Finally we present research directions that can potentially improve performance of social recommender systems including exploiting the heterogeneity of social networks and weak dependence connections, microcosmic investigation of users and items, considering temporal information in rating and social information, understanding the role of negative relations, and integrating cross-media data.

Accomplishment 7: "Exploiting Homophily Effect for Trust Prediction"(hTrust)

- **Research Problem Studied:** Trust plays a crucial role for online users who seek reliable information. However, in reality, user-specified trust relations are very sparse, i.e., a tiny number of pairs of users with trust relations are buried in a disproportionately large number of pairs without trust relations, making trust prediction a daunting task. As an important social concept, however, trust has received growing attention and interest. Social theories are developed for understanding trust. Homophily is one of the most important theories that explain why trust relations are established. Exploiting the homophily effect for trust prediction provides challenges and opportunities. In this paper, we embark on the challenges to investigate the trust prediction problem with the homophily effect. First, we delineate how it differs from existing approaches to trust prediction in an unsupervised setting. Next, we formulate the new trust prediction problem into an optimization problem integrated with homophily, empirically evaluate our approach on two datasets from real-world product review sites, and compare with representative algorithms to gain a deep understanding of the role of homophily in trust prediction.
- **Key Contributions:** In this paper, we study the problem of exploiting homophily effect for trust prediction. First we conduct experiments on datasets from real-world product review sites to demonstrate the existence of homophily in trust relations. Homophily regularization is then introduced to capture homophily effect in trust relations. An unsupervised framework is proposed, incorporating low-rank matrix factorization and homophily regularization. Extensive experiments are conducted to evaluate the proposed framework on real-world trust relation datasets and the experimental results demonstrate the effectiveness of our proposed framework as well as the role of homophily regularization for trust prediction.

Accomplishment 8: "Exploiting Local and Global Social Context for Recommendation"

- **Research Problem Studied:** With the fast development of social media, the information overload problem becomes increasingly severe and recommender systems play an important role in helping online users find relevant information by suggesting information of potential interests. Social activities for online users produce abundant social relations. Social relations provide an independent source for recommendation, presenting both opportunities and challenges for traditional recommender systems. Users are likely to seek suggestions from both their local friends and users with high global reputations, motivating us to exploit social relations from local and global perspectives for online recommender systems in this paper. We develop approaches to capture local and global social relations, and propose a novel framework LOCABAL taking advantage of both local and global social context for recommendation. Empirical results on real-world datasets demonstrate the effectiveness of our proposed framework and further experiments are conducted to understand how local and global social context work for the proposed framework.

- **Key Contributions:** The availability of social relations presents both challenges and opportunities for traditional recommender systems. In this paper, we investigate how to exploit local and global social context for recommendation. To capture local social context, we force that the user preferences of two socially connected users are correlated as suggested by social correlation theories and we also study the connections between our proposed approach and existing approaches. Ratings from users with high reputations are more likely to be trustworthy; therefore, to capture global social context, we use user reputation scores to weight the importance of their ratings. With these solutions, we propose a framework LOCABAL to integrate local and global social context for recommendation. Experimental results on real-world data sets show that the proposed framework LOCABAL outperforms representative social recommender systems. Further experiments are conducted to understand the working of LOCABAL.

Accomplishment 9: "A Tool for Collecting Provenance Data in Social Media"

- **Research Problem Studied:** In recent years, social media sites have provided a large amount of information. Recipients of such information need mechanisms to know more about the received information, including the provenance. Previous research has shown that some attributes related to the received information provide additional context, so that recipient can assess the amount of value, trust, and validity to be placed in the received information. Personal attributes of a user, including name, location, education, ethnicity, gender, and political and religious affiliations, can be found in social media sites. In this paper, we present a novel web-based tool for collecting the attributes of interest associated with a particular social media user related to the received information. This tool provides a way to combine different attributes available at different social media sites into a single user profile. Using different types of Twitter users, we also evaluate the performance of the tool in terms of number of attribute values collected, validity of these values, and total amount of retrieval time.
- **Key Contributions:** The provenance data collector tool aims to collect provenance attribute values of a user. By collecting such values of a user related to the received information, the tool could facilitate recipients to understand more about the received information. Data generated on social media sites is largely distributed and unstructured in nature. The proposed tool also provides a way to combine such distributed and unstructured social media data.

Accomplishment 10: "Recovering Information Recipients in Social Media via Provenance"

- **Research Problem Studied:** In recent years, social media has changed the way we interact and communicate. Although the existing structure of social media allows users to easily create, receive, and propagate pieces of information, many a time, users do not have background knowledge about the received information, including the provenance (sources or originators) of information, and other recipients who may have retransmitted

or modified the information. Providing such additional context to the received information can help users know how much value, trust, and validity should be placed in received information. To judge the credibility of the received piece of information, it is vital to know who are its sources, and how information propagates from sources to other social media users. In this paper, we are studying a novel research problem that facilitates a few known recipients to recover other unknown recipients, and seek the provenance of information. The experimental results with Facebook and Twitter datasets show that the proposed algorithm is effective in correctly recovering the unknown recipients and seeking the provenance of information.

- **Key Contributions:** Social media allows its users to share a vast amount of information with other users, but it provides no mechanism to know more about the received information for its users. In this paper, we aim to recover information recipients as well as seek the provenance by knowing a few nodes and using only link information in social networks. Information recipients exist along the paths from the sources to the known nodes. In this paper we seek the information propagation flow from the sources to the known nodes, and recover the most likely information recipients. Using the experiment results from the Facebook and Twitter datasets, we show that the proposed algorithm is effective in correctly recovering the information recipients and seeking the provenance of information.

Accomplishment 11: "Context-Aware Review Helpfulness Rating Prediction"

- **Research Problem Studied:** Online reviews play a vital role in the decision-making process for online users. Helpful reviews are usually buried in a large number of unhelpful reviews, and with the consistently increasing number of reviews, it becomes more and more difficult for online users to find helpful reviews. Therefore most online review websites allow online users to rate the helpfulness of a review and a global helpfulness score is computed for the review based on its available ratings. However, in reality, user-specified helpfulness ratings for reviews are very sparse a few reviews attract large numbers of helpfulness ratings while most reviews obtain few or even no helpfulness ratings. The available helpfulness ratings are too sparse for online users to assess the helpfulness of reviews. Also the helpfulness of a review is not necessarily equally useful for all users and users with different background may treat the helpfulness of a review very differently. The user idiosyncrasy of review helpfulness motivates us to study the problem of review helpfulness rating prediction in this paper. We first identify various types of context information, model them mathematically, and propose a context-aware review helpfulness rating prediction framework CAP. Experimental results demonstrate the effectiveness of the proposed framework and the importance of context awareness in solving the review helpfulness rating prediction problem.
- **Key Contributions:** In this paper we study the problem of review helpfulness rating prediction by exploiting context awareness to infer unknown helpfulness ratings

automatically, motivated by the fact that helpful reviews can be buried in large amounts of useless reviews and the user-specific helpfulness ratings are too sparse for online users to assess the helpfulness of reviews. We first show that the problem we study differs from review quality prediction problem and the item rating prediction problem. We extract four types of social context, i.e., author context, rater context, connection context and preference context, formulate them mathematically, and propose a context-aware helpfulness prediction framework CAP which exploits content context and various types of social context. Experimental results demonstrate that our proposed framework outperforms the state-of-the-art baseline methods with both cold-start and warm-start settings, and further experiments are conducted to understand the importance of context awareness in the proposed framework.

Accomplishment 12: "Seeking Provenance of Information in Social Media"

- **Research Problem Studied:** Social media has profoundly impacted the way people interact and communicate. Social media propagates breaking news and disinformation alike fast and on an unsurpassed scale. Because of its democratization nature, social media users can easily produce, receive and propagate a piece of information without necessarily providing traceable information. Thus, there are no means for a user to verify the provenance (also known as, sources or originators) of information. The disinformation can cause tragic consequences to society and individuals. This work aims to take advantage of characteristics of social media to provide a solution to the problem of lacking traceable information. Such knowledge can provide additional context to the received information such that a user can assess how much value, trust, and validity should be placed in received information. In this paper, we are studying a novel research problem that facilitates a few known recipients (less than 1% of the total recipients) to seek the provenance of information by recovering how it has own from its originators. The proposed methodology exploits easily computable node centralities of a large social media network. The experimental results with Facebook and Twitter datasets show that the proposed mechanism is effective in correctly identifying the additional recipients and seeking the provenance of information.
- **Key Contributions:** Social media allows its users to share vast amount of information with other users, but it lacks mechanisms that provide traceable knowledge about the received information for its users. In this paper, we study a novel research problem that facilitates a few P-nodes (less than 1% of total recipients) to seek the provenance of information by identifying how it has own from its originators. To this end, we first formally present the problem and provide the complexity analysis. Then, use the Facebook and Twitter datasets to show the existence of two hypotheses: Degree Propensity and Closeness Propensity. The proposed methodology then exploits these hypotheses to provide not only the critical information about the provenance, but also the most likely provenance paths. Finally using the experimental results with the Facebook

and Twitter datasets, we show that the proposed algorithm is effective in correctly identifying the additional transmitters, and seeking the provenance of information.

Accomplishment 13: "A Tool for Assisting Provenance Search in Social Media"

- **Research Problem Studied:** In recent years, social media sites are witnessing an information explosion. Determining the reliability of such a large amount of information is a major area of research. Information provenance (aka, sources or origin) provides a way to measure the reliability of information in social networks. The main challenge in seeking provenance is the availability of suitable data consisting of sufficient unique propagation paths. Current research on provenance in social media uses synthetically generated propagation paths. Although these proposed approaches are theoretically significant, it is still a challenge to apply and evaluate them on social media. Hence, knowledge of the actual propagation paths for a piece of information will be a valuable asset in provenance search. This paper presents a tool for capturing the propagation network of a given tweet or URL (Uniform Resource Locator) in the Twitter network. Researchers can use this tool to collect information propagation data, design effective strategies for determining the provenance, and gain information about the tweet such as impact, growth rate and users influencing the spread. An overview of the user interface and the architecture of the system is provided. Two case studies, one relating to disinformation in riot situations and another on corporate involvement in education has been presented to demonstrate the effectiveness of the system for seeking provenance information.
- **Key Contributions:** The paper presents a tool to obtain the spread of a given tweet or URL on the twitter network. The tool presents researchers with a propagation network to assist in seeking the provenance path of a given tweet. The provenance path gives additional information to assess the reliability of a given piece of data in social media.

Accomplishment 14: "Provenance Data in Social Media"

- **Book Overview:** Social media shatters the barrier to communicate anytime anywhere for people of all walks of life. The publicly available, virtually free information in social media poses a new challenge to consumers who have to discern whether a piece of information published in social media is reliable. For example, it can be difficult to understand the motivations behind a statement passed from one user to another, without knowing the person who originated the message. Additionally, false information can be propagated through social media, resulting in embarrassment or irreversible damages. Provenance data associated with a social media statement can help dispel rumors, clarify opinions, and confirm facts. However, provenance data about social media statements is not readily available to users today. Currently, providing this data to users requires changing the social media infrastructure or offering subscription services. Taking advantage of social media features, research in

this nascent field spearheads the search for a way to provide provenance data to social media users, thus leveraging social media itself by mining it for the provenance data. Searching for provenance data reveals an interesting problem space requiring the development and application of new metrics in order to provide meaningful provenance data to social media users. This lecture reviews the current research on information provenance, explores exciting research opportunities to address pressing needs, and shows how data mining can enable a social media user to make informed judgements about statements published in social media.

- Table of Contents: Information Provenance in Social Media / Provenance Attributes / Provenance via Network Information / Provenance Data

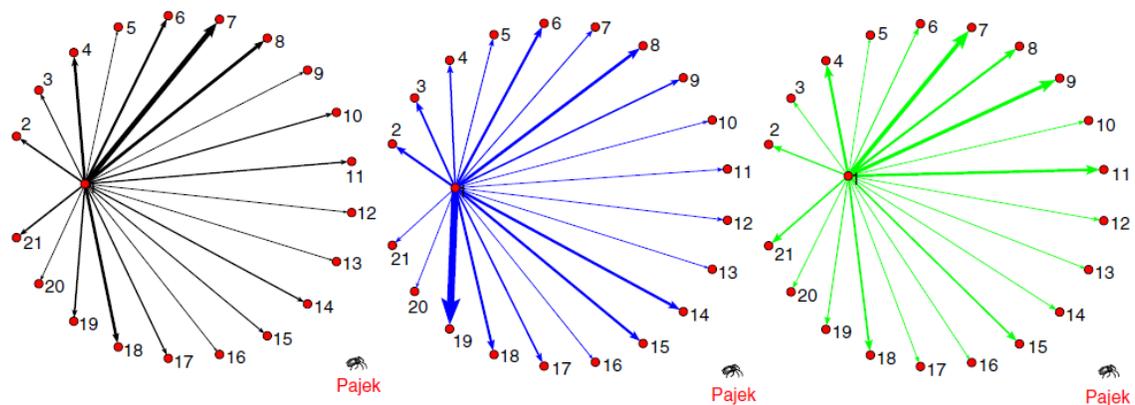
Accomplishment 15: "User Vulnerability and its Reduction on a Social Networking Site"

- Research Problem Studied: Privacy and security are major concerns for many users of social media. When users share information (e.g., data and photos) with friends, they can make their friends vulnerable to security and privacy breaches with dire consequences. With the continuous expansion of a user's social network, privacy settings alone are often inadequate to protect user's profile. In this research, we aim to address some critical issues related to privacy protection: (1) How can we measure and assess individual user's vulnerability? (2) With the diversity of one's social network friends, how can one figure out an effective approach to maintaining balance between vulnerability and social utility? In this work, first we present a novel way to define vulnerable friends from an individual user's perspective. User vulnerability is dependent on whether or not the user's friends' privacy settings protect the friend and the individual's network of friends (which includes the user). We show that it is feasible to measure and assess user vulnerability, and reduce one's vulnerability without changing the structure of a social networking site. The approach is to unfriend one's most vulnerable friends. However, when such a vulnerable friend is also socially important, unfriending him would significantly reduce one's own social status. We formulate this novel problem as vulnerability minimization with social utility constraints. We formally define the optimization problem, and provide an approximation algorithm with a proven bound. Finally, we conduct a large-scale evaluation of new framework using a Facebook dataset. We resort to experiments and observe how much vulnerability an individual user can decrease by unfriending a vulnerable friend. We compare performance of different unfriending strategies and discuss the security risk of new friend request. Additionally, by employing different forms of social utility, we confirm that balance between user vulnerability and social utility can be practically achieved.
- Key Contributions: We propose a feasible approach to a novel problem of identifying a user's vulnerable friends on a social networking site. Our work differs from existing work addressing social networking privacy by introducing a vulnerability-centered

approach to a user security on a social networking site. On most social networking sites, privacy related efforts have been concentrated on protecting individual attributes only. However, users are often vulnerable through community attributes. Unfriending vulnerable friends can help protect users against the security risks. Based on our study of over 2 million users, we find that users are either not careful or not aware of security and privacy concerns of their friends. Our model clearly highlights the impact of each new friend on a user’s privacy. Our approach does not require the structural change of a social networking site and aims to maximally reduce a user’s vulnerability while minimizing his social utility loss. The work formulates a novel problem of constrained vulnerability reduction suggests a feasible approach, and demonstrates that the problem of constrained vulnerability reduction is solvable.

Accomplishment 16: “mTrust: Discerning Multi-Faceted Trust in a Connected World”

- **Research Problem Studied:** The issue of trust has attracted increasing attention from the community of social media research. Trust, as a social concept, naturally has multiple facets, indicating multiple and heterogeneous trust relationships between users. Here is a multifaceted trust example from Epinions. Figure 1(a) shows single trust relationships between user 1 and his 20 friends. Here, we can see that user 7 is the more trustable for user 1. Figures 1(b) and 1(c) show their multifaceted trust relationships in the categories “home and garden” and “restaurants” respectively. For the category “home and garden” user 7 is not necessary the most trusted friend of user 1. This shows that trust relationships in different categories vary. Thus, people trust others differently in different facets.



(a) Single Trust (b) Trust in home and garden (c) Trust in restaurants

Figure 1: Single trust and multifaceted trust relationships of one use in Epinions.

(Note: The thickness of a line indicates the level of trust.)

There are two challenges to study in obtaining multifaceted trust between users: first, the representation of multiple and heterogeneous trust relationships between users, and second, estimating the strength of multifaceted trust. Traditionally, trust is represented by

an adjacency matrix. However, this cannot capture the multifaceted trust relations. We developed a new algorithm, mTrust, which extends a matrix representation to a tensor representation, adding an extra dimension for facet description. Previous work observed a strong correlation between trust and user similarity in the context of rating systems. Therefore, it is reasonable to embed trust strength inference in rating prediction. Thus, to evaluate the usefulness of multifaceted trust, this work embeds the multifaceted trust inference in the framework of rating prediction.

- **Key Contributions:** Interesting findings from the experiments are that (1) more than 20% of reciprocal links are heterogeneous, (2) more than 14% transitive trust relations are heterogeneous, and (3) more than 11% of cocitation trust relations are heterogeneous. With these findings, mTrust can be applied to many online tasks such as improving rating prediction, enabling facet-sensitive ranking, and making status theory applicable to reciprocal links.

Accomplishment 17: “eTrust: Understanding Trust Evolution in an Online World”

- **Research Problem Studied:** Most existing research about online trust assumes static trust relations between users. As we are informed by social sciences, trust evolves as humans interact. Little work exists studying trust evolution in an online world. Researching online trust evolution faces unique challenges because more often than not, available data is from passive observation. In this paper, we leverage social science theories to develop a methodology that enables the study of online trust evolution. In particular, we propose a framework of evolution trust, eTrust, which exploits the dynamics of user preferences in the context of online product review. We present technical details about modeling trust evolution, and perform experiments to show how the exploitation of trust evolution can help improve the performance of online applications such as rating and trust prediction.
- **Key Contributions:** We study online trust evolution in the context of product review. By exploiting the correlation between user preferences and trust relations, we propose a framework, eTrust, to understand the evolution of trust in an online world and apply eTrust to various online applications such as rating prediction and trust prediction. Interesting findings are observed in our experiments using real-world data, Epinions; and eTrust can be applied to improve the performance of rating prediction and trust prediction.

Accomplishment 18: “Minimizing User Vulnerability and Retaining Social Utility in Social Media”

- **Research Problem Studied:** Privacy and security are major concerns for many users of social media. When users share information (e.g., data and photos) with friends, they can make their friends vulnerable to security and privacy breach with dire consequences. In our earlier work, we show that it is feasible to measure user vulnerability and reduce one's vulnerability without changing the structure of a social networking site. The

approach is to unfriend one's most vulnerable friends. However, when such a vulnerable friend is also socially important, unfriending him would significantly reduce one's own social status. In this work, we address the problem of vulnerability minimization with minimum social utility losses. This work extends the existing vulnerability reduction model to a more general form. Using a general model, we formulate the two discrete optimization problems. Both problems are NP-hard.

- **Key Contributions:** We formally formulate the optimization problem, propose an approximation algorithm with a proven bound, and conduct empirical experiments with different forms of social utility on a large-scale Facebook dataset for performance evaluation and comparison. Our work differs from existing work addressing social networking privacy. Our approach does not require the structural change of a social networking site and aims to maximally reduce a user's vulnerability while minimizing his social utility loss.